


[> home](#) [> about](#) [> feedback](#) [> login](#)

U.S. Patent &amp; Trademark Office



Try the *new* Portal  
design

Give us your opinion after using  
it.

## Search Results

Search Results for: **[raid<AND>((david a. patterson<AND>((patterson) )) )]**  
Found **48** of **125,779** searched.

## Search within Results

[> Advanced Search](#)[> Search Help/Tips](#)

Sort by: Title Publication Publication Date Score Binder

Results 1 - 20 of 48 short listing



1

2

3



- 1** Electronic structure of materials using self-interaction corrected density functional theory 84%



Adolfy Hoisie , Stefan Goedecker , Jurg Hutter

**Proceedings of the 1996 ACM/IEEE conference on Supercomputing (CDROM)**

November 1996

We have developed an highly efficient electronic structure code, for parallel computers using message passing. The algorithm takes advantage of the natural parallelism in quantum chemistry problems to obtain very high performance even on a large number of processors. Most of the terms which scale cubically with respect to the number of atoms have been eliminated allowing the treatment of very large systems. It uses one of the most precise versions of Density Functional Theory, namely Self-I ...

- 2** The HP AutoRAID hierarchical storage system 84%



J. Wilkes , R. Golding , C. Staelin , T. Sullivan

**ACM SIGOPS Operating Systems Review , Proceedings of the fifteenth ACM symposium on Operating systems principles** December 1995

Volume 29 Issue 5

- 3** RAID-II: a high-bandwidth network file server 82%



A. L. Drapeau , K. W. Shirriff , J. H. Hartman , E. L. Miller , S. Seshan , R. H. Katz , K. Lutz , D. A. Patterson , E. K. Lee , P. M. Chen , G. A. Gibson

**ACM SIGARCH Computer Architecture News , Proceedings of the 21ST annual international symposium on Computer architecture** April 1994

Volume 22 Issue 2

In 1989, the RAID (Redundant Arrays of Inexpensive Disks) group at U. C. Berkeley built a prototype disk array called RAID-I. The bandwidth delivered to clients by RAID-I was severely limited by the memory system bandwidth of the disk array's host workstation. We designed our second prototype, RAID-H, to deliver more of the disk array bandwidth to file server clients. A custom-built crossbar memory system called the XBUS board connects the disks directly to the high-speed network, allowing data ...

- 4 An analytic performance model of disk arrays 82%  
Edward K. Lee , Randy H. Katz  
**ACM SIGMETRICS Performance Evaluation Review , Proceedings of the 1993 ACM SIGMETRICS conference on Measurement and modeling of computer systems** June 1993  
Volume 21 Issue 1
- 5 The TickerTAIP parallel RAID architecture 82%  
Pei Cao , Swee Boon Lim , Shivakumar Venkataraman , John Wilkes  
**ACM SIGARCH Computer Architecture News , Proceedings of the 20th annual international symposium on Computer architecture** May 1993  
Volume 21 Issue 2  
Traditional disk arrays have a centralized architecture, with a single controller through which all requests flow. Such a controller is a single point of failure, and its performance limits the maximum size that the array can grow to. We describe here TickerTAIP, a parallel architecture for disk arrays that distributed the controller functions across several loosely-coupled processors. The result is better scalability, fault tolerance, and flexibility. This paper presents the Tic ...
- 6 Performance of a disk array prototype 82%  
Ann L. Chervenak , Randy H. Katz  
**ACM SIGMETRICS Performance Evaluation Review , Proceedings of the 1991 ACM SIGMETRICS conference on Measurement and modeling of computer systems** April 1991  
Volume 19 Issue 1
- 7 Performance consequences of parity placement in disk arrays 82%  
Edward K. Lee , Randy H. Katz  
**ACM SIGARCH Computer Architecture News , Proceedings of the fourth international conference on Architectural support for programming languages and operating systems** April 1991  
Volume 19 Issue 2
- 8 Database session 4: heterogeneous and distributed systems: A reliable storage management layer for distributed information retrieval systems 80%  
Charles L. A. Clarke , Philip L. Tilker , Allen Quoc-Luan Tran , Kevin Harris , Antonio S. Cheng  
**Proceedings of the twelfth international conference on Information and knowledge management** November 2003  
We present a storage management layer that facilitates the implementation of parallel information retrieval systems, and related applications, on networks of workstations. The storage management layer automates the process of adding and removing nodes, and implements a dispersed mirroring strategy to improve reliability. When nodes are added and removed, the document collection managed by the system is redistributed for load balancing purposes. The use of dispersed mirroring minimizes the impact ...
- 9 An evaluation of redundant arrays of disks using an Amdahl 5890 80%  
Peter M. Chen , Garth A. Gibson , Randy H. Katz , David A. Patterson  
**ACM SIGMETRICS Performance Evaluation Review , Proceedings of the 1990 ACM SIGMETRICS conference on Measurement and modeling of computer systems** April 1990  
Volume 18 Issue 1  
Recently we presented several disk array architectures designed to increase the data rate and I/O rate of supercomputing applications, transaction processing, and file systems [Patterson 88]. In this paper we present a hardware performance measurement of two of these architectures, mirroring and rotated parity. We see how throughput for these two architectures is affected by response time requirements, request sizes, and read to write

ratios. We find that for applications with large ...

# 10 Beyond striping: the bridge multiprocessor file system 80%



P. C. Dibble , M. L. Scott

**ACM SIGARCH Computer Architecture News** September 1989

Volume 17 Issue 5

High-performance parallel computers require high-performance file systems. Exotic I/O hardware will be of little use if file system software runs on a single processor of a many-processor machine. We believe that cost-effective I/O for large multiprocessors can best be obtained by spreading both data and file system computation over a large number of processors and disks. To assess the effectiveness of this approach, we have implemented a prototype system called Bridge, and have studied its perf ...

# 11 Analysis of methods for scheduling low priority disk drive tasks 80%



Eitan Bachmat , Jiri Schindler

**ACM SIGMETRICS Performance Evaluation Review , Proceedings of the 2002 ACM SIGMETRICS international conference on Measurement and modeling of computer systems** June 2002

Volume 30 Issue 1

This paper analyzes various algorithms for scheduling low priority disk drive tasks. The derived closed form solution is applicable to class of greedy algorithms that include a variety of background disk scanning applications. By paying close attention to many characteristics of modern disk drives, the analytical solutions achieve very high accuracy---the difference between the predicted response times and the measurements on two different disks is only 3% for all but one examined workload. This ...

# 12 Specifying data availability in multi-device file systems 80%



John Wilkes , Raymie Stata

**Proceedings of the 4th workshop on ACM SIGOPS European workshop** September 1990

# 13 Maximizing performance in a striped disk array 80%



Peter M. Chen , David A. Patterson

**ACM SIGARCH Computer Architecture News , Proceedings of the 17th annual international symposium on Computer Architecture** May 1990

Volume 18 Issue 3

Improvements in disk speeds have not kept up with improvements in processor and memory speeds. One way to correct the resulting speed mismatch is to stripe data across many disks. In this paper, we address how to stripe data to get maximum performance from the disks. Specifically, we examine how to choose the striping unit, i.e. the amount of logically contiguous data on each disk. We synthesize rules for determining the best striping unit for a given range of workloads.

# 14 Optimizing throughput in a workstation-based network file system over a 80%




high bandwidth local area network

Theodore Faber

**ACM SIGOPS Operating Systems Review** January 1998

Volume 32 Issue 1

This paper describes methods of optimizing a client/server network file system to advantage of high bandwidth local area networks in a conventional distributed computing environment. The environment contains hardware that removes network and disk bandwidth bottlenecks. The remaining bottlenecks at clients include excessive context switching, inefficient data translation, and cumbersome data encapsulation methods. When these are removed, the null-write performance of a current implementation of S ...

**15** Multiprocessor out-of-core FFTs with distributed memory and parallel disks 80% (extended abstract)

Thomas H. Cormen , Jake Wegmann , David M. Nicol


**Proceedings of the fifth workshop on I/O in parallel and distributed systems**

November 1997

**16** Petal: distributed virtual disks 80% Edward K. Lee , Chandramohan A. Thekkath**Proceedings of the seventh international conference on Architectural support for programming languages and operating systems** September 1996


Volume 31 , 30 Issue 9 , 5

The ideal storage system is globally accessible, always available, provides unlimited performance and capacity for a large number of clients, and requires no management. This paper describes the design, implementation, and performance of Petal, a system that attempts to approximate this ideal in practice through a novel combination of features. Petal consists of a collection of network-connected servers that cooperatively manage a pool of physical disks. To a Petal client, this collection appear ...

**17** Serverless network file systems 80% Thomas E. Anderson , Michael D. Dahlin , Jeanna M. Neefe , David A. Patterson , Drew S. Roselli , Randolph Y. Wang**ACM Transactions on Computer Systems (TOCS)** February 1996

Volume 14 Issue 1

We propose a new paradigm for network file system design: serverless network file systems. While traditional network file systems rely on a central server machine, a serverless system utilizes workstations cooperating as peers to provide all file system services. Any machine in the system can store, cache, or control any block of data. Our approach uses this location independence, in combination with fast local area networks, to provide better performance and scalability th ...

**18** Striping in a RAID level 5 disk array 80% Peter M. Chen , Edward K. Lee**ACM SIGMETRICS Performance Evaluation Review , Proceedings of the 1995 ACM SIGMETRICS joint international conference on Measurement and modeling of computer systems** May 1995

Volume 23 Issue 1

Redundant disk arrays are an increasingly popular way to improve I/O system performance. Past research has studied how to stripe data in non-redundant (RAID Level 0) disk arrays, but none has yet been done on how to stripe data in redundant disk arrays such as RAID Level 5, or on how the choice of striping unit varies with the number of disks. Using synthetic workloads, we derive simple design rules for striping data in RAID Level 5 disk arrays given varying amounts of workload information. We t ...

**19** A new approach to I/O performance evaluation: self-scaling I/O 80% benchmarks, predicted I/O performance

Peter M. Chen , David A. Patterson

**ACM Transactions on Computer Systems (TOCS)** November 1994

Volume 12 Issue 4

Current I/O benchmarks suffer from several chronic problems: they quickly become obsolete; they do not stress the I/O system; and they do not help much in understanding I/O system performance. We propose a new approach to I/O performance analysis. First, we propose a self-scaling benchmark that dynamically adjusts aspects of its workload according to the performance characteristic of the system being measured. By doing so, the benchmark automatically scales across current and future systems ...

**20 RAID: high-performance, reliable secondary storage**

80%



Peter M. Chen , Edward K. Lee , Garth A. Gibson , Randy H. Katz , David A. Patterson

**ACM Computing Surveys (CSUR)** June 1994

Volume 26 Issue 2

Disk arrays were proposed in the 1980s as a way to use parallelism between multiple disks to improve aggregate I/O performance. Today they appear in the product lines of most major computer manufacturers. This article gives a comprehensive overview of disk arrays and provides a framework in which to organize current and future work. First, the article introduces disk technology and reviews the driving forces that have popularized disk arrays: performance and reliability. It discusses the tw ...

---

**Results 1 - 20 of 48****short listing****1****2****3**

---

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2004 ACM, Inc.